UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/677,661	10/02/2003	Joel Howard Schopp	AUS920030739US1	9970

35525 7590 03/02/2007

IBM CORP (YA)
C/O YEE & ASSOCIATES PC
P.O. BOX 802333
DALLAS, TX 75380

EXAMINER	
WALTER, CRAIG E	
ART UNIT	**PAPER NUMBER**
2188	

SHORTENED STATUTORY PERIOD OF RESPONSE	MAIL DATE	DELIVERY MODE
2 MONTHS	03/02/2007	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

If NO period for reply is specified above, the maximum statutory period will apply and will expire 6 MONTHS from the mailing date of this communication.

UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

MAILED

MAR 0 2 2007

Technology Center 2100

BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

Application Number: 10/677,661
Filing Date: October 02, 2003
Appellant(s): SCHOPP, JOEL HOWARD

Theodore D. Fay III
For Appellant

EXAMINER'S ANSWER

This is in response to the appeal brief filed 5 December 2006 appealing from the Office action mailed 18 August 2006.

(1) Real Party in Interest

A statement identifying by name the real party in interest is contained in the brief.

(2) Related Appeals and Interferences

The examiner is not aware of any related appeals, interferences, or judicial proceedings which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

(3) Status of Claims

The statement of the status of claims contained in the brief is correct.

(4) Status of Amendments After Final

No amendment after final has been filed.

(5) Summary of Claimed Subject Matter

The summary of claimed subject matter contained in the brief is correct.

(6) Grounds of Rejection to be Reviewed on Appeal

The appellant's statement of the grounds of rejection to be reviewed on appeal is correct.

(7) Claims Appendix

The copy of the appealed claims contained in the Appendix to the brief is correct.

(8) Evidence Relied Upon

US PG Publication 2003/0065886	Olarig et al.	04-2003
US Patent 6,212,605	Arimilli et al.	04-2001
US Patent 6,581,115	Arimilli et al.	06-2003
US PG Publication 2003/0084372	Mock et al.	05-2003
US Patent 5,784,702	Greenstein et al.	07-1998

(9) Grounds of Rejection

The following ground(s) of rejection are applicable to the appealed claims:

Claim Rejections - 35 USC § 102

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

> A person shall be entitled to a patent unless -
> (e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

Claims 1, 4, 8, 10 and 13 are rejected under 35 U.S.C. 102(e) as being anticipated by Olarig et al. (US PG Publication 2003/0065886 A1), hereinafter Olarig.

As for claims 1, 8 and 10, Olarig teaches a method (and system) for assigning processors to partitions in a multi-processor data processing system, the method comprising:

generating optimal allocation sets for unallocated processors in the multi-processor data processing system for a cache level, wherein each optimal allocation set includes an allocation of unallocated processors to at least one partition (the cache can be partitioned variety of different manners in order to optimize the processors' utilization of the cache memory. For example, each of the eight processors may have their own unique partition of the cache if needed to optimize allocation - paragraph 0019, all lines. Note the processors are unallocated until a determination is made as to which processors are to be allocated to their respective cache partitions);

determining whether a first set in the optimal allocation sets matches requirements for a set of partitions selected for the data processing system (cache optimization can be determined based on cache utilization - paragraph 0021, all lines. In other words, if a running hit average of the cache hits a certain value; the partitions can be allocated to processors in order to optimize the system's efficiency. The system checks to see if the parameters are met or matched to make the determination); and

responsive to a match existing, removing processors in the first set from the unallocated processors to form a group of removed processors, wherein

cache usage by the group of removed processors is optimized for the cache level (paragraph 0021, all lines - again the cache will be repartitioned based on the statistics previously calculated for system efficiency. It is also worthy to note that the system is not just limited in resizing the previous partitions, but it is also capable of reallocating the cache partitions based on the statistical information. This way, the system can assign processors to partitions in order to optimize the system for that cache level - see paragraphs 0026 and 0027, all lines).

Note also with respect to claim 8, Olarig teaches the additional step of selecting the highest cache level that has been unprocessed in the multi-processor system (as per the example discussed in the previously cited sections, Olarig discuss determining the new partition allocation of the cache). In this example, the cache itself is the highest level of memory (i.e. L1).

Additionally, Olarig teaches generating a different set of allocations based on match requirements if no match occurs (again the partitioning will occur based on the need of the system as determined by the statistics collected by the system (paragraphs 0019 and 0020, all lines)).

As for claims 4 and 13, Olarig teaches removing the partition for the set of partitions in response to allocating the first set to the partition (paragraph 0019, all lines - once allocated, the partition is no longer available to the remaining processors until the reallocation is determined with the statistical information collected by the system).

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

Claims 3 and 12 are rejected under 35 U.S.C. 103(a) as being unpatentable over Olarig as applied to claims 1 and 10 above, and in further view of Arimilli et al. (US Patent 6,212,605 B1, hereinafter Arimilli '605.

As for claim 3 and 12, though Olarig teaches his system as including multiple cache levels, he fails to teach repeating the steps of claim 1 (and claim 10) for each cache level by moving from the highest cache level, to each successively lower cache level.

Arimilli '605 however teaches a method for controlling eviction of cache blocks to override the eviction of a value, which is reserved for a later operation. More specifically, Arimilli '605 teaches performing system operations on a hierarchical, multi-level cache (i.e. setting an address reservation flag), by starting at the highest cache level (in this case, L2), and then repeating the system operation on the next lower level - col. 3, lines 46-65.

It would have been obvious to one of ordinary skill in the art at the time of the invention for Olarig to further include Arimilli's ('605) system of eviction override into his

own system of cache partitioning, in order to facilitate repeating the partitioning/allocation techniques of his memory system on each successive level of cache. By doing so, Olarig would benefit by having a means of improving the overall processor performance of his system by avoiding unnecessary evictions to the cache. This way the overall performance will improve, by avoiding a scenario in which a block that has be unnecessarily evicted must be reloaded into the cache as taught by Arimilli '605 - col. 7, lines 38-58.

Claims 6-7 and 15-16 rejected under 35 U.S.C. 103(a) as being unpatentable over Olarig as applied to claims 1 and 10 above, and in further view of Arimilli et al., hereinafter Arimilli '115 (US Patent 6,581,115 B1).

As for claims 6 and 15, though Olarig fails to teach his system as a symmetric multi-processor system (SMP) including L3 and L2 cache, Arimilli '115 teaches a processing system with configurable memory bus and scalability port which specifically apply to an SMP (see abstract), that includes both L2 and L3 cache (col. 5, lines 46-55 and col. 6, lines 19-34).

As for claims 7 and 16, Olarig fails to teach the group of removed processors as being located on multi-chip modules, wherein optimal allocations of the group of removed processors are based on the location of the processors on the modules.

Arimilli '115 however teaches a system containing SMPs, which are located on multi-chip modules (col. 2, lines 20-42). Arimilli '115 further teaches allocating and connecting the processors on the modules to the cache based on bandwidth utilization.

In other words, smaller number of processors can be allocated to particular memory partition if the bandwidth allocation needs to be lower. This way, the memory usage can be optimized based on the location of the modules (i.e. more processors located on a bus that has higher allocation, and less processors located on buses with lower bandwidth allocation) -col. 6, lines 19-35.

It would have been obvious to one of ordinary skill in the art at the time of the invention for Olarig to further include Arimilli's ('115) data processing system to his own system of cache partitioning. By doing so, Olarig would be able to dynamically configure a bus allocation of a processor chip utilizing hardware and software allocation methods, and efficiently allocate processor chip bus bandwidth in order to improve the overall data processing system as taught by Arimilli '115 (col. 3, lines 44-56).

Claims 9, 17 and 20 rejected under 35 U.S.C. 103(a) as being unpatentable over Olarig, and in further view of Mock et al. (US PG Publication 2003/0084372 A1), hereinafter Mock.

As for claims 9 and 17, Olarig fails to teach a memory for storing instructions, which the processors use to carry out the remaining steps of the claim (the remaining steps are discussed under the rejection of claim 1 presented *supra*). Mock however teaches an apparatus for data recovery optimization in a logically partitioned computer system, in which the processors in the multiprocessor system use instructions stored in main memory to carry out their respective functions (paragraph 0034, all lines 1-11).

As for claim 20, Olarig teaches removing the partition for the set of partitions in response to allocating a set to the partition (paragraph 0019, all lines - once allocated the partition is no longer available to the remaining processors until the reallocation is determined with the statistical information collected by the system).

It would have been obvious to one of ordinary skill in the art at the time of the invention for Olarig to further include Mock's apparatus for data recovery in a logically partitioned computer system. By doing so, Olarig would benefit by not only having a means of recovering data in his system, but have a system that is capable of implementing a recovery strategy that accurately reflects the current system configuration. This would in turn increase overall system performance, as discussed by Mock in paragraph 0014, all lines.

Claim 19 is rejected under 35 U.S.C. 103(a) as being unpatentable over the combined teachings of Olarig and Mock as applied to claim 17 above, and in further view of Arimilli '605.

As for claim 19, though Olarig teaches his system as including multiple cache levels, he fails to teach repeating the steps of claim 1 for each cache level by moving from the highest cache level, to each successively lower cache level.

Arimilli '605 however teaches a method for controlling eviction of cache blocks to override the eviction of a value, which is reserved for a later operation. More specifically, Arimilli '605 teaches performing system operations on a hierarchical, multi-level cache (i.e. setting an address reservation flag), by starting at the highest cache

level (in this case, L2), and then repeating the system operation on the next lower level - col. 3, lines 46-65.

It would have been obvious to one of ordinary skill in the art at the time of the invention for Olarig to further include Arimilli's ('605) system of eviction override into his own system of cache partitioning, in order to facilitate repeating the partitioning/allocation techniques of his memory system on each successive level of cache. By doing so, Olarig would benefit by having a means of improving the overall processor performance of his system by avoiding unnecessary evictions to the cache. This way the overall performance will improve by avoiding a scenario in which it is necessary to reload a block that been unnecessarily evicted from the cache memory as taught by Arimilli '605 - col. 7, lines 38-58.

Claims 21-23 rejected under 35 U.S.C. 103(a) as being unpatentable over the combined teachings of Olarig and Mock as applied to claim 17 above, and in further view of Arimilli '115.

As for claims 21-22, though Olarig fails to teach his system as a symmetric multi-processor system (SMP) including L3 and L2 cache, Arimilli '115 teaches a processing system with configurable memory bus and scalability port which specifically applies to an SMP (see abstract), that includes both L2 and L3 cache (col. 5, lines 46-55 and col. 6, lines 19-34).

As for claim 23, Olarig fails to teach the processors as being located on multi-chip modules, wherein optimal allocations are based on the location of the processors on the modules.

Arimilli '115 however teaches a system containing SMPs, which are located on multi-chip modules (col. 2, lines 20-42). Arimilli '115 further teaches allocating and connecting the processors on the modules to the cache based on bandwidth utilization. In other words, smaller number of processors can be allocated to particular memory if the bandwidth allocation needs to be lower. This was, the memory usage can be optimized based on the location of the modules (i.e. more processors located on a bus that has higher allocation, and less processors located on buses with lower bandwidth allocation) -col. 6, lines 19-35.

It would have been obvious to one of ordinary skill in the art at the time of the invention for Olarig to further include Arimilli's ('115) data processing system to his own system of cache partitioning. By doing so, Olarig would be able to dynamically configure a bus allocation of a processor chip utilizing hardware and software allocation methods, and efficiently allocate processor chip bus bandwidth in order to improve the overall data processing system as taught by Arimilli '115 (col. 3, lines 44-56).

Claims 24-25, and 27 are rejected under 35 U.S.C. 103(a) as being unpatentable over Olarig as applied to claims 1 and 10 above, and in further view of Greenstein et al. (US Patent 5,784,702), hereinafter Greenstein.

As for claim 24, though Olarig teaches all the limitations of claim 1, he fails to specifically teach the partitions as comprising logical partitioned data processing systems, wherein the at least one partition comprise at least one logical partitioned data processing system, and wherein the set of partitions comprise a set of logical partitioned data processing systems.

Greenstein however teaches a system and method for dynamically performing resource reconfiguration in a logically partitioned data processing system in which the system's resources are logically partitioned without the need for operator involvement (col. 2, lines 36-59). Fig. 3 illustrates one example of Greenstein's system configuration prior to the reconfiguration process. Specifically, the configuration comprises, *inter alia*, two partitions (elements 301, and 302), multiple storage elements and CPUs. Additionally, the configuration is subdivided into three logical partitions (A, B and C as illustrated in the figure) - col. 8 line 30, through col. 9, line 6. Figure 4 depicts the partitioning after reconfiguration takes place. It is now apparent from Fig. 4 that each partition comprises at least one data processing system, and the set of partitions comprise a set of partitioned data processing systems.

It would have been obvious to one of ordinary skill in the art at the time of the invention for Olarig to further include Greenstein's system for dynamically performing resource reconfiguration in a logically partitioned data processing system into his own system for dynamic cache partitioning. By doing so, Olarig could exploit the benefits of more efficiently partitioning hardware resources via a dynamic process without the need and/or aid of operator involvement, which in turn would save a considerable amount of

down time by minimizing/eliminating the steps of deactivating all logical partitions, performing re-initialization, then reactivating the logical partitions during the reconfiguration process as taught by Greenstein in col. 1, lines 60-67 and col. 2, lines 35-59.

Claims 25 and 27 are rejected for similar reasons to claim 24.

Claims 26 and 28 are rejected under 35 U.S.C. 103(a) as being unpatentable over the combined teachings of Olarig and Mock as applied to claims 9 and 17 above, and in further view of Greenstein.

Though the combined teachings of Olarig and Mock meet all of the limitations of claims 9 and 17, they fail to teach the remaining limitations in clams 26 and 28. Greenstein however teaches these limitations per the rejection of claim 24 above.

It would have been obvious to one of ordinary skill in the art at the time of the invention for Olarig to further include Greenstein's system for dynamically performing resource reconfiguration in a logically partitioned data processing system into his own system for dynamic cache partitioning. By doing so, Olarig could exploit the benefits of more efficiently partitioning hardware resources via a dynamic process without the need and/or aid of operator involvement, which in turn would save a considerable amount of down time by minimizing/eliminating the steps of deactivating all logical partitions, performing re-initialization, then reactivating the logical partitions during the reconfiguration process as taught by Greenstein in col. 1, lines 60-67 and col. 2, lines 35-59.

(10) Response to Argument

Appellant's arguments have been fully considered but they are not persuasive.

Under heading A. Grounds of rejection 1 (claims 1, 4, 8, 10 and 13), Appellant questions whether Olarig fails to anticipate claims 1, 4, 8, 10 and 13, and sets forth three specific arguments as discussed below.

Under heading A.1. Response to Rejection, Appellant contends, "each and every feature of the presently claimed invention is not identically shown in the cited reference, arranged as they are in the claims". More specifically, Appellant contends, "Olarig never actually states or discloses, "wherein each optimal allocation set includes an allocation of unallocated processors to at least one partition,"". This argument is not persuasive. Referring again to paragraph 0019, all lines of Olarig, cache memory is partitioned, and dedicated to a different processor. This is done in order to optimize the cache utilization of Olarig's system (paragraph 0010, all lines). Appellant is further attributing an overall narrow definition of "allocate" to this claim limitation. Based on Examiner's broadest reasonable interpretation consistent with Applicant's specification (MPEP § 2111), an allocation is taken to mean assignment for a particular purpose. This definition is not only consistent with Applicant's specification, but also consistent within its meaning to one of ordinary skill in the art. It is worthy to note that Olarig specifically teaches allocating previously unallocated processors to partitions of a particular cache level, because he uses termination such as "allocate" and "dedicate", rather than "re-allocate", or "re-dedicate". One of ordinary skill in the art understands the former terms to mean dedicate or allocating something that has no previous

allocation or dedication, rather than the latter terms which are taken to mean allocate or dedicate something that has already been dedicated or assigned previously. Also note Olarig's teachings are directed to a *dynamic partitioning* (paragraphs 0010 through 0011, all lines), meaning the system can allocate previously unallocated resources, rather than being limited to a static allocation where no modifications to the resources may be performed. Since Olarig teaches allocating or dedicating the processors to cache partitions that previously did not exist (i.e. prior to the dynamic partitioning and assignment), it is well known in the art that these resources (i.e. cache and processors) were previously unallocated.

In order to further illustrate this point, Examiner has provided a simplified version of Olarig's Fig. 1 at the end of this section. This drawing will hereinafter be referred to as "Examiner's Fig. 1". Examiner's Fig. 1 illustrates Olarig's eight processors and one cache prior to allocation (i.e. the top portion of the figure), and further depicts his eight processors and one cache after allocation (i.e. bottom portion). Before allocation, each of the eight processors are coupled to the cache, however they are not allocated (i.e. unallocated) with respect to any particular partition of the cache since the cache has yet to be partitioned. Olarig teaches in paragraph 0019, all lines, by way of example, a dynamic partitioning scheme, which splits the cache into eight unique partitions, and assigns each processor a partition in order to optimize this level of the cache (as depicted in the bottom portion of Examiner's Fig. 1). Once this partitioning takes place, each partition of the cache is uniquely assigned to a processor by a one-one correspondence, hence each processor is uniquely assigned to at least one partition of

the cache as required by the instant claim. Further, the group of eight processors originally are *unallocated* with respect any particular partition, however once the partitioning takes place they are no longer unallocated (i.e. removed), and are now classified as a group of allocated processors with respect to cache partitioning assignment. It is worthy to note that this partitioning occurs *dynamically* (i.e. changing the state from unallocated to allocated, rather than being statically assigned to processors as a separate embodiment of the invention) to optimize the cache level, as evidenced by Olarig's teachings in paragraphs 0010 through 0011, all lines.

Appellant further alleges that Examiner relied upon what Olarig has broadly taught, rather that what Olarig discloses (i.e. "[t]hus, even if Olarig "broadly taught" the claimed feature of "an allocation of unallocated processors," an assertion that is clearly wrong, Olarig would still not anticipate claim 1 because Olarig does not explicitly disclose this claimed feature"). This argument however is not persuasive, as Examiner maintains that Olarig discloses allocating unallocated processors in paragraph 0019, all lines as he teaches taken previously unassigned/undedicated/unallocated processors and assigns them to particular partitions of a cache. Please refer again to Examiner's Fig. 1, *infra*.

Appellant additionally contends that "Olarig does not necessarily teach, "wherein each optimal set includes an allocation of unallocated processors to a least one partition" as claimed". This argument however is not persuasive, as Examiner maintains that Olarig discloses allocating unallocated processors in paragraph 0019, all lines, as he teaches taking previously unassigned/undedicated/unallocated processors,

and assigning them to particular partitions of a cache. This allocation is performed in order to reach an arrangement, which is optimized for the processor to cache partition assignment (paragraph 0010, all lines).

Lastly, Appellant contends there is "no basis or reason to assume the existence of "an allocation of unallocated processors to at least one partition" of caches". This statement is based on Appellant's contention that "the examiner simply assume[ed] the existence of unallocated processors". This argument however is not persuasive. Examiner has not assumed the existence of unallocated processors. Again, it is worthy to note that Olarig is specifically teaches allocating previously unallocated processors to partitions of a particular cache level, because he uses termination such as "allocate" and "dedicate", rather than "re-allocate", or "re-dedicate". One of ordinary skill in the art understands the former terms to mean dedicate or allocating something that has no previous allocation or dedication, rather than the latter terms which are taken to mean allocate or dedicate something that has already been dedicated or assigned previously. Since Olarig teaches allocating or dedicating the processors to cache partitions, it is well known in the art that these resources (i.e. cache and processors) were previously unallocated. Please refer again to Examiner's Fig. 1, *infra*

Appellant continues by stating that Examiner's application of Olarig's teachings to claim 1 is vis-à-vis backwards. More specifically, Appellant contends that Olarig partitions caches and then assigns the partitions of caches to processors, whereas by contrast, claim 1 requires an allocation of unallocated processors to at least one partition. Examiner however maintains Olarig's disclosure pertaining to partitioning the

cache, and dedicating (i.e. assigning) processors to a respective partition anticipates this limitation - paragraph 0019, all lines. In other words, the example in which each of the eight processors are assigned to a unique partition clearly demonstrates a one-one correspondence of each processor being allocated to a particular partition of the cache.

Appellant further contends that Olarig fails to teach "generating optimal allocation sets". More specifically, Appellant asserts that the example taught by Olarig by assigning eight cache bins to eight processors is by no means optimal, and additionally argues that even if it is assumed that the allocation is optimized, the sets are for unallocated cache partitions and not for unallocated processors. This argument however is not persuasive. Examiner maintains that allocation as taught by Olarig is in fact generated for optimal allocation. Referring to paragraphs 0019 through 0020, all lines, Olarig describes a non-exhaustive list of possible allocations of processors and PCI resources. Once the allocation is complete, the resources (i.e. processor) have exclusive rights to its respective partition. In other words, Olarig's disclosure allows for the optimal utilization of the cache by allocating the available resources in such a fashion as to minimize contention and/or data thrashing. Additionally note paragraph 0021, all lines specifically addresses cache partitioning and reallocation of partitioning as being a means for optimization. As for Appellant's argument that optimal allocation is for cache partitions and not for unallocated processors (by referring again to paragraphs 0019 through 0020, all lines), Olarig if fact teaches partitioning the cache based on the resources available (i.e. PCI resources and processors), hence once the allocation is complete, the unallocated processors are now allocated as a an optimal

set. Paragraph 0010, all lines further evidences the fact that Olarig's disclosure aims to optimize memory allocated to processors.

Appellant further concedes that Olarig teaches optimization of cache partitions, however contends that Olarig fails to teach "determining whether a first set in the optimal allocation set matches requirements for a set of partitions selected for the data processing system". This argument however is not persuasive. Examiner maintains that optimal set includes allocated processors per the first limitation of the claim, "each allocation set includes an allocation of unallocated processor to a least one partition". In other words, said set does in fact include allocated processors despite Appellant's contention otherwise. The specific arrangement and allocation of processors to cache is dynamically performed for the specific purpose of cache/system optimization (paragraphs 0010 through 0011, all lines).

Lastly Appellant contends that Olarig does not teach "responsive to a match existing, removing processors in the first set from the unallocated processors to form a group of removed processors, wherein cache usage by the group of removed processors is optimized for the cache level". By contrast Appellant asserts that Olarig teaches optimizing cache partitions by reallocating the sizes of the cache partitions relative to the workload of the cache partitions, and (at best) teaches removing caches from unallocated caches. This argument however is not persuasive. Examiner maintains once a processor is allocated it is no longer unallocated, hence it is removed from the group of unallocated processors. The process of gathering statistical data on the utilization of the cache is specifically tailored for the purpose of optimizing the

utilization of the cache partition at that particular level of the cache. Please refer again to Examiner's Fig. 1, *infra*.

Under heading A.2. Rebuttal to Examiner's Response, Appellant contends, "examiner has provided no basis or foundation for the assertion that assigning a partition to a processors is equivalent to assigning a processor to a partition. The examiner has the burden to establish this fact". This argument however is not persuasive. Referring again to paragraph 0019, all lines of Olarig, each partition of the cache is assigned to a unique processor. One of ordinary skill in the art would understand that once a cache is assigned to partition, the processor is now assigned to the cache (most certainly this is true with respect to a one-one correspondence example for eight processors as described by Olarig). Olarig teaches an example of partitioning the first level of a cache into eight unique bins, and subsequently assigning each of these partitions to a unique processor. Once this assignment transpires, each processor is uniquely assigned to a cache partition; hence Olarig anticipates the limitation "allocation of unallocated processors to at least one partition" as recited in this claim.

Appellant further contends, "simply because the examiner states that the claim language is reflexive does not mean that reversing the order of the claim language results in the same claim or that reversing the order of assignment results in the same invention". This argument however is not persuasive. The claims requires, *inter alia*, "generating optimal allocation sets for unallocated processors in the multi-processor data processing system for a cache level, wherein each optimal allocation set includes

an allocation of unallocated processors to at least one partition". Examiner maintains that Olarig anticipates this limitation. More specifically, once the cache partitions are each assigned to a unique processor, the system has in fact generated an optimized allocation set of processors to at least one partition of a level one cache, as taught in paragraph 0019, all lines. Examiner further maintains the assertion that assigning a partition to a processor is equivalent to assigning a processor to a partition in view of the scope of the instant claim, in that they both result in "generating optimal allocation sets for unallocated processors in the multi-processor data processing system for a cache level, wherein each optimal allocation set includes an allocation of unallocated processors to at least one partition" as required by the claim. To reiterate, the example provided where eight partitions are assigned to eight processors in Olarig's disclose results in eight processors that are now uniquely assigned to respective cache partitions as required by the instant claim.

Appellant further contends, "the claim language does not recite generating optimal allocation sets by assigning processors to cache partitions. Instead, the claim recites generating optimal allocation sets... for a cache level...". Appellant further states, "[t]his feature is not equivalent to assigning processors to cache partitions. This argument however is not persuasive. Examiner maintains that the claim requires "generating optimal allocation sets for unallocated processors in the multi-processor data processing system for **a cache level**, wherein each optimal allocation set includes an allocation of unallocated processors to at least one partition" (emphasis added). Olarig discusses leveled cache in paragraph 0007, all lines. Though he does not

explicitly indicate which level of cache is being partitioned and assigned to unique processors in paragraph 0019, all lines, one of ordinary skill in the art would understand any cache to be a first level cache, hence Olarig anticipates generating the optimal allocation set for a cache level (i.e. L1) as required by the claim. To reiterate, in paragraph 0019, all lines, Olarig in fact teaches partitioning one cache (in this example a level one cache) for optimization by uniquely assigning partitions from that one cache to multiple processors.

Appellant further asserts "Olarig teaches that the cache itself is optimized", however "In stark contrast, claim 1, requires that, "each optimal allocation set includes an allocation of unallocated processors to at least one partition."". Lastly, Appellant asserts "If a processor uses an optimized cache in Olarig, that processor is not part of the group of removed processor and that processor is not taught by Olarig to be in an allocation of unallocated processors, as claimed", and that "no unallocated processors are allocated as an optimal set in Olarig". This argument however is not persuasive. Examiner maintains that an optimal set of previously unallocated processors is generated once Olarig partitions the cache, and assigns processors to each partition. This is in fact the purpose for Olarig to dynamically partition the cache. This optimal set of allocations is directed to, and aimed at optimized the system's utilization of the cache – Olarig, paragraph 0019, all lines.

Appellant lastly assets, "The examiner's assertions are clearly wrong. As shown above, Olarig never explicitly discloses "an allocation" of "unallocated processors," as claimed. This argument however is not persuasive. Examiner maintains that once the

cache is partitioned and assigned to unique processors, the system has now taken processors, which were not previously assigned to partitions (i.e. unallocated), and assigns them (i.e. allocates) to a unique partition within the cache level. Please refer again to Examiner's Fig. 1, *infra*.

Under heading A.3. Summary of why the Examiner's rebuttal is insufficient to establish that Olarig anticipates claim 1 or the other claims in the grouping of claims, Appellant contends "[t]he examiner also assumes that in Olarig processors are removed from the imaginary allocation of unallocated processors because once a processor is assigned cache partition is assigned to a cache partition that the processor is removed from the conceptual set of processors that have not yet been allocated. However, again, Olarig does not actually teach this feature". This argument is not persuasive, as Examiner maintains that once processors are allocated or assigned to partitions they are no longer unallocated, and before the cache is partitioned they are not allocated to partitions, hence they are unallocated with respect to a cache partition, and subsequently allocated to partitions once the assignment transpires.

Under heading B. Grounds of rejection 2 (claims 3 and 12), Appellant contends that Examiner has failed to state a *prima facie* obviousness rejection against claims 3 and 12 under 35 U.S.C. § 103(a) over Olarig in view of Arimilli '605. Appellant sets forth three specific arguments as discussed below.

Under the heading B.1. The proposed combination does not teach or suggest all of the features of the claims, Appellant argues Olarig fails to teach all of the elements of

claims 1 and 10, and therefore fails to each all of the elements of claim 3, which directly depends on claim 1. Appellant further contends that Arimilli '605 fails to each or suggest any of the features of claim 1, therefore neither Olarig nor Arimilli '605, alone or in combination, teach all of the features of claim 1 and claim 3. This argument however is not persuasive as Examiner maintains that Olarig teaches each and every element of claim 1 per the arguments and rejection presented *supra*.

Under the heading B.2. No motivation exists to combine Olarig and Arimilli '605 because the address different problems, Applicant additionally alleges that no motivation exists to combine Olarig and Arimilli '605 because they address different problems. More specifically, Appellant contends that one of ordinary skill in the art would not combine the reference to achieve the invention to claim 3 because the references are directed toward solving different problems. Appellant contends that since Olarig is allegedly directed to solving the problems of reducing cache thrashing, and that Arimilli '605 is allegedly directed to the problem of speeding up implementation of lwarx/stwcx semantics, these references are addressing completely distinct problems that are unrelated to each other. This argument however is not persuasive, as Examiner maintains that both Olarig and Arimilli '605 are in fact directed to improving the efficiency of a multi-processor system. More specifically, Olarig is directed to a more efficient cache, which utilizes dynamic cache partitioning for a plurality of cache entries. The cache partitioning improves the cache efficiency through targeted replacement - paragraph 0011, all lines. Arimilli '605 is directed to improving the efficiency of a multi-processor system, capable of more efficiently handing cache

reservations and evictions (i.e. replacement) of cache lines - col. 5, lines 42-52. Olarig and Arimilli '605 are both therefore directed to helping improving the efficiency of a multi-processor data system, and further improving the means in which data is written to, and read from a cache. Therefore, both Olarig and Arimilli '605 are directed to solving similar problems within the meaning of 35 USC § 103(a). As stated in the previously within the rejection above, It would have been obvious to one of ordinary skill in the art at the time of the invention for Olarig to further include Arimilli's ('605) system of eviction override into his own system of cache partitioning, in order to facilitate repeating the partitioning/allocation techniques of his memory system on each successive level of cache. By doing so, Olarig would benefit by having a means of improving the overall processor performance of his system by avoiding unnecessary evictions to the cache. This way the overall performance will improve, by avoiding a scenario in which a block that has be unnecessarily evicted must be reloaded into the cache as taught by Arimilli '605 - col. 7, lines 38-58.

Under the heading B.3. The examiner has failed to state a prima facie case obviousness rejection because Arimilli '605 is non-analogous art, Appellant contends that Arimilli '605 is not in the same field of endeavor of claim 3, and Arimilli '605 is not reasonably pertinent to the particular problem with which Applicant's were concerned. More specifically, Appellant contends that Arimilli '605 is not in the same field of endeavor of claim 3 because Arimilli '605 is in the field of handling cache reservations and evictions in a multi-processor data processing system. Appellant additionally

contends that Arimilli '605 is not reasonably pertinent to the particular problem with which Applicants were concerned.

Examiner however maintains that Arimilli '605 is in fact both within the same field of Applicant's endeavor, and reasonably pertinent to the particular problem with which the inventor was concerned. More specifically, Arimilli '605 is directed to improving the efficiency of cache within a multi-processor system - col. 5, lines 42-52. Appellant's field of endeavor, and particular problem to solve is related to managing and optimizing a cache within a multi-processor environment (i.e. as evidenced by Appellant's title "Cache optimized logical partitioning of a symmetric multi-processor data processing system"), and further evidenced by Appellant's description of the technical field of invention as disclosed on page 1, lines 1-7 of Appellant's original specification. Olarig, Arimilli '605, and Appellant field of endeavor are therefore directed to helping improve the efficiency of a multi-processor system, and further directed to improving the means in which data is written to, and read from a cache. Arimilli '605 is therefore analogous art within the meaning of 35 USC § 103(a).

Under heading C. Grounds of rejection 3 (claims 6, 7, 15 and 16), Appellant contends that Examiner has failed to state a *prima facie* obviousness rejection against claims 6, 7, 15, and 16 under 35 U.S.C. § 103(a) over Olarig in view of Arimilli '115. Appellant sets forth three specific arguments as discussed below.

Under the heading C.1. The proposed combination does not teach or suggest all of the features of the claims, Appellant argues Olarig fails to teach all of the elements of

claim 1, and therefore fails to each all of the elements of claim 6, which directly depends on claim 1. Appellant further contends that Arimilli '115 fails to each or suggest any of the features of claim 1, therefore neither Olarig nor Arimilli '115, alone or in combination, teach all of the features of claim 1 and claim 6. This argument however is not persuasive as Examiner maintains that Olarig teaches each and every element of claim 1 per the arguments and rejection presented *supra*.

Under the heading C.2. No motivation exists to combine the references because the address different problems, Appellant alleges that no motivation exists to combine Olarig and Arimilli '115 because they address different problems. More specifically, Appellant contends that one of ordinary skill in the art would not combine the references to achieve the invention to claim 6 because the references are directed toward solving different problems. Appellant contends that since Olarig is allegedly directed to solving the problems of reducing cache thrashing, and that Arimilli '115 is allegedly directed to the problem of managing buses of a processor chip between competing external components to increase overall efficiency of the data processing system, these references are addressing completely distinct problems that are unrelated to each other. This argument however is not persuasive, as Examiner maintains that both Olarig and Arimilli '115 are in fact directed to improving the overall efficiency of multi-processing systems. More specifically, Olarig is directed to a more efficient multi-processing data system, which utilizes dynamic cache partitioning for a plurality of cache entries. The cache partitioning improves the cache efficiency through targeted replacement - paragraph 0011, all lines. Arimilli '115 is directed to improving the

efficiency of a multi-processor system - col. 3, lines 64-67. Olarig and Arimilli '115 are both therefore directed to improving the efficiencies of critical resources (i.e. bus allocation and memory management) within a multi-processor system, and are therefore directed to solving similar problems within the meaning of 35 USC § 103(a). As stated previously within the rejection *supra*, It would have been obvious to one of ordinary skill in the art at the time of the invention for Olarig to further include Arimilli's ('115) data processing system to his own system of cache partitioning. By doing so, Olarig would be able to dynamically configure a bus allocation of a processor chip utilizing hardware and software allocation methods, and efficiently allocate processor chip bus bandwidth in order to improve the overall data processing system as taught by Arimilli '115 (col. 3, lines 44-56).

Under the heading C.3. The examiner has failed to state a prima facie case obviousness rejection because Arimilli '115 is non-analogous art, Appellant contends that Arimilli '115 is not in the same field of endeavor of claim 6, and Arimilli '115 is not reasonably pertinent to the particular problem with which Applicant's were concerned. More specifically, Appellant contends that Arimilli '115 is not in the same field of endeavor of claim 6 because Arimilli '115 is in the field of managing buses of a processor chip between competing external components to increase overall efficiency of the data processing system. Appellant additionally contends that Arimilli '115 is not reasonably pertinent to the particular problem with which Applicants were concerned.

Examiner however maintains that Arimilli '115 is in fact both within the same field of Applicant's endeavor, and reasonably pertinent to the particular problem with which

the inventor was concerned. More specifically, Arimilli '115 is directed to improving the efficiency of a multi-processor system - col. 3, lines 64-67. Appellant's field of endeavor, and particular problem to solve is related to managing and optimizing a critical resource (i.e. cache) within a multi-processor environment (i.e. as evidenced by Appellant's title "Cache optimized logical partitioning of a symmetric multi-processor data processing system"), and further evidenced by Appellant's description of the technical field of invention as disclosed on page 1, lines 1-7 of Appellant's original specification. As stated above, both Appellant and Arimilli '115 are both therefore directed to improving the efficiencies of critical resources (i.e. bus allocation and memory management) within a multi-processor system. Olarig and Arimilli '115 are both therefore directed to improving the efficiencies of critical resources (i.e. bus allocation and memory management) within a multi-processor system, and are therefore analogous art within the meaning of 35 USC § 103(a).

Under heading D. Grounds of rejection 4 (claims 9, 17 and 20), Appellant contends that Examiner has failed to state a *prima facie* obviousness rejection against claims 9, 17 and 02 under 35 U.S.C. § 103(a) over Olarig in view of Mock. Appellant sets forth three specific arguments as discussed below.

Under the heading D.1. The proposed combination does not teach or suggest all of the features of the claims, Appellant argues Olarig in view of Mock fail to teach all of the elements of claim 9, based on a similar rationale as applied to claim 1, *supra*. Appellant further contends that Mock fails to each or suggest any of the features of claim 9, therefore neither Olarig nor Mock, alone or in combination, teach all of the

features of claim 9. This argument however is not persuasive as Examiner maintains that Olarig in view of Mock teach each and every element of claim 9 per the arguments and rejection presented *supra*, and that Olarig teaches each and every claim limitation of claim 1 as per the rejection and arguments presented *supra*.

Under the heading D.2. No motivation exists to combine Olarig and Mock because the address different problems, Appellant contends that one of ordinary skill in the art would not combine the reference to achieve the invention to claim 9 because the references are directed toward solving different problems. Appellant contends that since Olarig is allegedly directed to solving the problems of reducing cache thrashing, and that Mock is allegedly directed to the problem of recovery of data in a logically partitioned environment, these references are addressing completely distinct problems that are unrelated to each other. This argument however is not persuasive, as Examiner maintains that both Olarig and Mock are in fact directed to optimizing a multi-processor system. More specifically, Olarig is directed to a more efficient cache, which utilizes dynamic cache partitioning for a plurality of cache entries. The cache partitioning improves the cache efficiency through targeted replacement - paragraph 0011, all lines. This partitioning occurs within a multi-processor system. Mock is directed to a system capable of logically partitioning system resources for multi-processors in order to mitigating system inefficiency (i.e. unnecessary I/O operation) - paragraphs 0017 through 0019, all lines. Olarig and Mock are both therefore directed to helping improve multi-processor system efficiencies, and are therefore directed to solving similar problems within the meaning of 35 USC § 103(a). As stated in the

previously within the rejection above, It would have been obvious to one of ordinary skill in the art at the time of the invention for Olarig to further include Mock's apparatus for data recovery in a logically partitioned computer system. By doing so, Olarig would benefit by not only having a means of recovering data in his system, but have a system that is capable of implementing a recovery strategy that accurately reflects the current system configuration. This would in turn increase overall system performance, as discussed by Mock in paragraph 0014, all lines.

Under the heading D.3. The examiner has failed to state a prima facie case obviousness rejection because Mock is non-analogous art, Appellant contends that Mock is not in the same field of endeavor of claim 9, and Mock is not reasonably pertinent to the particular problem with which Applicant's were concerned. More specifically, Appellant contends that Mock is not in the same field of endeavor of claim 9 because Mock is in the field recovery of certain data in a digital computer system. Appellant additionally contends that Mock is not reasonably pertinent to the particular problem with which Applicants were concerned.

Examiner however maintains that Mock is in fact both within the same field of Applicant's endeavor, and reasonably pertinent to the particular problem with which the inventor was concerned. More specifically, Mock is directed to a system capable of logically partitioning system resources for multi-processors in order to mitigate system inefficiency (i.e. unnecessary I/O operation) - paragraphs 0017 through 0019, all lines. Appellant's field of endeavor, and particular problem to solve is related to managing and optimizing a cache within a multi-processor environment (i.e. as evidenced by

Appellant's title "Cache optimized logical partitioning of a symmetric multi-processor data processing system"), and further evidenced by Appellant's description of the technical field of invention as disclosed on page 1, lines 1-7 of Appellant's original specification. Olarig and Mock are both therefore directed to helping improve multi-processor system efficiencies, and are therefore analogous art within the meaning of 35 USC § 103(a).

Under heading E. Grounds of rejection 5 (claim 19), Appellant contends that Examiner has failed to state a *prima facie* obviousness rejection against claim 19 under 35 U.S.C. § 103(a) over Olarig in view Mock in view of Arimilli '605. Appellant sets forth three specific arguments as discussed below.

Under the heading E.1. The proposed combination does not teach or suggest all of the features of the claims, Appellant argues Olarig fails to teach all of the elements of claim 1 common with claim 19. Appellant further contends that neither Mock nor Arimilli '605 teach or suggest any of the features of claim 19, therefore neither Olarig, nor Mock, nor Arimilli '605, alone or in combination, teach all of the features of claim 19. This argument however is not persuasive as Examiner maintains that Olarig teaches each and every element of claim 1 (and therefore the features in common with claim 19) per the arguments and rejection presented *supra*.

Under the heading E.2. No motivation exists to combine the references because the address different problems, Appellant alleges that no motivation exists to combine Olarig, Mock and Arimilli '605 because they address different problems. More

specifically, Appellant contends that one of ordinary skill in the art would not combine the reference to achieve the invention to claim 19 because the references are directed toward solving different problems. Appellant contends that no common sense reason exists to establish that one of ordinary skill would reasonably be expected to look for a solution to the problem facing the invention. This argument however is not persuasive, as Examiner maintains that Olarig, Mock and Arimilli '605 are in fact directed to improving multi-processor systems. As stated previously, Olarig, Mock and Arimilli '605 are directed to improving the efficiency of critical resources within multi-processor systems, and are therefore directed to solving similar problems within the meaning of 35 USC § 103(a). As stated in the previously within the rejection above, It would have been obvious to one of ordinary skill in the art at the time of the invention for Olarig to further include Arimilli's ('605) system of eviction override into his own system of cache partitioning, in order to facilitate repeating the partitioning/allocation techniques of his memory system on each successive level of cache. By doing so, Olarig would benefit by having a means of improving the overall processor performance of his system by avoiding unnecessary evictions to the cache. This way the overall performance will improve by avoiding a scenario in which it is necessary to reload a block that been unnecessarily evicted from the cache memory as taught by Arimilli '605 - col. 7, lines 38-58.

Under the heading E.3. The examiner has failed to state a prima facie case obviousness rejection because Mock and Arimilli '605 are non-analogous art, Appellant contends that Mock and Arimilli '605 are not in the same field of endeavor of claim 19,

and both Mock and Arimilli '605 are not reasonably pertinent to the particular problem with which Applicant's were concerned.

Examiner however maintains that both Mock and Arimilli '605 are in fact both within the same field of Applicant's endeavor, and reasonably pertinent to the particular problem with which the inventor was concerned. More specifically, Mock, Arimilli '605, and Applicant's instant invention are directed to improving the efficiency of critical system resources in multiprocessing systems; therefore they are analogous art within the meaning of 35 USC § 103(a).

Under heading F. Grounds of rejection 6 (claims 21-23), Appellant contends that Examiner has failed to state a *prima facie* obviousness rejection against claim 21-23 under 35 U.S.C. § 103(a) over Olarig in view Mock in view of Arimilli '115. Appellant sets forth three specific arguments as discussed below.

Under the heading F.1. The proposed combination does not teach or suggest all of the features of the claims, Appellant argues Olarig fails to teach all of the elements of claim 1 common with claim 21. Appellant further contends that neither Mock nor Arimilli '115 teach or suggest any of the features of claim 21, therefore neither Olarig, nor Mock, nor Arimilli '115, alone or in combination, teach all of the features of claim 21. This argument however is not persuasive as Examiner maintains that Olarig teaches each and every element of claim 1 (and therefore the features in common with claim 21) per the arguments and rejection presented *supra*.

Under the heading F.2. No motivation exists to combine the references because the address different problems, Appellant alleges that no motivation exists to combine Olarig, Mock and Arimilli '115 because they address different problems. More specifically, Appellant contends that one of ordinary skill in the art would not combine the reference to achieve the invention to claim 21 because the references are directed toward solving different problems. Appellant contends that no common sense reason exists to establish that one of ordinary skill would reasonably be expected to look for a solution to the problem facing the invention. This argument however is not persuasive, as Examiner maintains that Olarig, Mock and Arimilli '115 are in fact directed to improving multi-processor systems. As stated previously, Olarig, Mock and Arimilli '115 are directed to improving the efficiency of critical resources within multi-processor systems, and are therefore directed to solving similar problems within the meaning of 35 USC § 103(a). As stated in the previously within the rejection above, It would have been obvious to one of ordinary skill in the art at the time of the invention for Olarig to further include Arimilli's ('115) data processing system to his own system of cache partitioning. By doing so, Olarig would be able to dynamically configure a bus allocation of a processor chip utilizing hardware and software allocation methods, and efficiently allocate processor chip bus bandwidth in order to improve the overall data processing system as taught by Arimilli '115 (col. 3, lines 44-56).

Under the heading F.3. The examiner has failed to state a prima facie case obviousness rejection because Mock and Arimilli '115 are non-analogous art, Appellant contends that Mock and Arimilli '115 are not in the same field of endeavor of claim 21,

and both Mock and Arimilli '115 are not reasonably pertinent to the particular problem with which Applicant's were concerned.

Examiner however maintains that both Mock and Arimilli '115 are in fact both within the same field of Applicant's endeavor, and reasonably pertinent to the particular problem with which the inventor was concerned. More specifically, Mock, Arimilli '115, and Applicant's instant invention are directed to improving the efficiency of critical system resources in multiprocessing systems; therefore they are analogous art within the meaning of 35 USC § 103(a).

Under heading G. Grounds of rejection 7 (claims 24, 25 and 27), Appellant contends that Examiner has failed to state a *prima facie* obviousness rejection against claims 24, 25 and 27 under 35 U.S.C. § 103(a) over Olarig in view of Greenstein. Appellant sets forth three specific arguments as discussed below.

Under the heading G.1. The proposed combination does not teach or suggest all of the features of the claims, Appellant argues Olarig fails to teach all of the elements of claim 1, and therefore fails to each all of the elements of claim 24, which directly depends on claim 1. Appellant further contends that Greenstein fails to each or suggest any of the features of claim 1, therefore neither Olarig nor Greenstein, alone or in combination, teach all of the features of claim 1 and claim 24. This argument however is not persuasive as Examiner maintains that Olarig teaches each and every element of claim 1 per the arguments and rejection presented *supra*.

Under the heading G.2. No motivation exists to combine Olarig and Greenstein because the address different problems, Appellant contends that one of ordinary skill in the art would not combine the reference to achieve the invention to claim 24 because the references are directed toward solving different problems. Appellant contends that since Olarig is allegedly directed to solving the problems of reducing cache thrashing, and that Greenstein is allegedly directed to the problem of physically partitioning logical partitioned systems without disconnecting power to the physical systems, these references are addressing completely distinct problems that are unrelated to each other. This argument however is not persuasive, as Examiner maintains that both Olarig and Greenstein are in fact directed optimizing critical resources in a multiprocessor system. More specifically, Olarig is directed to a more efficient cache, which utilizes dynamic cache partitioning for a plurality of cache entries. The cache partitioning improves the cache efficiency through targeted replacement in a multi-processor system - paragraph 0011, all lines. Greenstein is directed configuring and reconfiguring system resources in a partitioned environment to improve system efficiency in a multi-processor system (col. 2, lines 36-58 and Fig. 3). Olarig and Greenstein are both therefore directed to helping improve the means in which data is written to, and read from a cache, and are therefore directed to solving similar problems within the meaning of 35 USC § 103(a). As stated in the previously within the rejection above, It would have been obvious to one of ordinary skill in the art at the time of the invention for Olarig to further include Greenstein's system for dynamically performing resource reconfiguration in a logically partitioned data processing system into his own

system for dynamic cache partitioning. By doing so, Olarig could exploit the benefits of more efficiently partitioning hardware resources via a dynamic process without the need and/or aid of operator involvement, which in turn would save a considerable amount of down time by minimizing/eliminating the steps of deactivating all logical partitions, performing re-initialization, then reactivating the logical partitions during the reconfiguration process as taught by Greenstein in col. 1, lines 60-67 and col. 2, lines 35-59.

Under the heading G.3. The examiner has failed to state a prima facie case obviousness rejection because Greenstein is non-analogous art, Appellant contends that Greenstein is not in the same field of endeavor of claim 24, and Greenstein is not reasonably pertinent to the particular problem with which Applicant's were concerned. More specifically, Appellant contends that Greenstein is not in the same field of endeavor of claim 24 because Greenstein is in the field of dynamic partition/merging support for an LPAR-mode processor. Appellant additionally contends that Greenstein is not reasonably pertinent to the particular problem with which Applicants were concerned.

Examiner however maintains that Greenstein is in fact both within the same field of Applicant's endeavor, and reasonably pertinent to the particular problem with which the inventor was concerned. More specifically, Greenstein is directed configuring and reconfiguring system resources in a partitioned environment to improve system efficiency in a multi-processor system (col. 2, lines 36-58 and Fig. 3). Appellant's field of endeavor, and particular problem to solve is related to managing and optimizing a

cache within a multi-processor environment (i.e. as evidenced by Appellant's title "Cache optimized logical partitioning of a symmetric multi-processor data processing system"), and further evidenced by Appellant's description of the technical field of invention as disclosed on page 1, lines 1-7 of Appellant's original specification. Olarig and Greenstein are both therefore directed to improving the efficiencies of critical system resources in a multi-processor system, and are therefore analogous art within the meaning of 35 USC § 103(a).

Under heading H. Grounds of rejection 8 (claims 26 and 28), Appellant contends that Examiner has failed to state a *prima facie* obviousness rejection against claim 26 and 28 under 35 U.S.C. § 103(a) over Olarig in view Mock in view of Greenstein.

Under the heading H.1. The proposed combination does not teach or suggest all of the features of the claims, Appellant argues Olarig in view of Mock fail to teach all of the elements of claim 26 for the reasons as stated with regards to claim 9. Appellant further contends that neither Mock nor Greenstein teach or suggest any of the features of claim 26, therefore neither Olarig, nor Mock, nor Greenstein, alone or in combination, teach all of the features of claim 26. This argument however is not persuasive as Examiner maintains that Olarig in view of Mock teach each and every element of claim 9 (and therefore the relevant portions of claim 26) per the arguments and rejection presented *supra*.

Under the heading H.2. No motivation exists to combine Olarig , Mock and Greenstein because the address different problems, Appellant contends that one of

ordinary skill in the art would not combine the reference to achieve the invention to claim 26 because the references are directed toward solving different problems. This argument however is not persuasive, as Examiner maintains that Olarig, Mock and Greenstein are in fact directed to improving multi-processor systems. As stated previously, Olarig, Mock and Greenstein are directed to improving the efficiency of critical resources within multi-processor systems, and are therefore directed to solving similar problems within the meaning of 35 USC § 103(a). As stated in the previously within the rejection above, it would have been obvious to one of ordinary skill in the art at the time of the invention for Olarig to further include Greenstein's system for dynamically performing resource reconfiguration in a logically partitioned data processing system into his own system for dynamic cache partitioning. By doing so, Olarig could exploit the benefits of more efficiently partitioning hardware resources via a dynamic process without the need and/or aid of operator involvement, which in turn would save a considerable amount of down time by minimizing/eliminating the steps of deactivating all logical partitions, performing re-initialization, then reactivating the logical partitions during the reconfiguration process as taught by Greenstein in col. 1, lines 60-67 and col. 2, lines 35-59.

Under the heading H.3. The examiner has failed to state a prima facie case obviousness rejection because Greenstein is non-analogous art, Appellant contends that Mock and Greenstein are not in the same field of endeavor of claim 26, and both Mock and Greenstein are not reasonably pertinent to the particular problem with which Applicant's were concerned.

Examiner however maintains that both Mock and Greenstein are in fact both within the same field of Applicant's endeavor, and reasonably pertinent to the particular problem with which the inventor was concerned. More specifically, Mock, Greenstein, and Applicant's instant invention are directed to improving the efficiency of critical system resources in multiprocessing systems; therefore they are analogous art within the meaning of 35 USC § 103(a).

Under the heading I. Common reasons why the examiner failed to state prima facie obviousness rejections against any of the claims, Appellant sets forth three specific arguments as discussed below.

Under heading I.1. None of the asserted reasons to combine the reference are proper teachings, suggestions, or motivations to combine the references, Appellant contends that none of the asserted reasons to combine the references are proper teachings, suggestions, or motivations to combine the references. This argument however is not persuasive. More specifically, Examiner maintains each and every 103(a) rejection set forth meets the four Graham v. John Deere Co. factual inquires set forth in MPEP § 2141, hence Examiner has met the burden of providing a *prima facie* case of obviousness. In each 103(a) rejection, Examiner has identified motivations that are **explicitly taught** by the secondary references. Examiner maintains that each of the explicitly cited motivations were cited to demonstrate why in fact one of ordinary skill in the art at the time of the invention would be motivated to combine the teachings. This practice is consistent with Title 35 of the Code, and further consistent with proper Office

practice and procedure as being sufficient to demonstrate a *prima facie* case of obviousness within the meaning of the statute.

Under heading I.2. The Examiner used impermissible hindsight when fashioning the rejections, Appellant asserts, "[t]he proposed combinations all rely on an interpretations of Olarig instead of on the explicit disclosures in Olarig". This argument however is not persuasive. Examiner maintains that Olarig discloses allocating unallocated processors in paragraph 0019, all lines as he teaches taken previously assigned/undedicated/unallocated processors and assigns them to particular partitions of a cache. Please refer again to Examiner's Fig. 1, *infra*.

Appellant additionally contends, "[t]he examiner used four secondary references in seven rejections in order to fashion all of the obviousness rejections. Each obviousness rejection only address one or a few claims". This argument however is not persuasive as Examiner fails to see the relevance as to how the number of references, and number of rejections applied bears any relevance on the patentability of the instant claims. Examiner maintains each and every 103(a) rejection set forth meets the four Graham v. John Deere Co. factual inquires set forth in MPEP § 2141, hence Examiner has met the burden of providing a *prima facie* case of obviousness.

Appellant additionally contends, "[n]ot once does the examiner actually offer a motivation to combine any of the references. Instead, the examiner only offers purported advantages to combines the references. This argument however is not persuasive. In each 103(a) rejection, Examiner has identified motivations that are

explicitly taught by the secondary references. Examiner maintains that each of the explicitly cited motivations were cited to demonstrate why in fact one of ordinary skill in the art at the time of the invention would be motivated to combine the teachings. This practice is consistent with Title 35 of the Code, and further consistent with proper Office practice and procedure as being sufficient to demonstrate a *prima facie* case of obviousness within the meaning of the statute.

Appellant additionally reiterates the previous contention that each of the references address different a problem. This argument however is not persuasive, as similarities were explicitly addressed for each combination of prior art per the arguments discussed *supra*.

Appellant lastly alleges, "the examiner only offer hypothetical advantages to combine the references instead of actual motivations to combine the references, then in light of the other facts the examiner again must have used personal opinion and impressible hindsight when combines the references. This argument is not persuasive, as Examiner has identified motivations that are **explicitly taught** by the secondary references, rather than hypothetical as alleged by Appellant. Examiner maintains that each of the explicitly cited motivations were cited to demonstrate why in fact one of ordinary skill in the art at the time of the invention would be motivated to combine the teachings. This practice is consistent with Title 25 of the Code, and further consistent with proper Office practice and procedure as being sufficient to demonstrate a *prima facie* case of obviousness within the meaning of the statute.

Under heading I.3. Rebuttal to Examiner's response, Appellant asserts, the examiner's proposed reasons why one of ordinary skill would combine the references all rely on hypothetical advantages to combine the references." This argument is not persuasive, as Examiner has identified motivations that are **explicitly taught** by the secondary references, rather than hypothetical as alleged by Appellant. Examiner maintains that each of the explicitly cited motivations were cited to demonstrate why in fact one of ordinary skill in the art at the time of the invention would be motivated to combine the teachings. This practice is consistent with Title 35 of the Code, and further consistent with proper Office practice and procedure as being sufficient to demonstrate a *prima facie* case of obviousness within the meaning of the statute.

In summary, Arguments under the headings A through I are not persuasive per the reasons discussed *supra*, and for those reasons, it is believed that the rejections should be sustained.

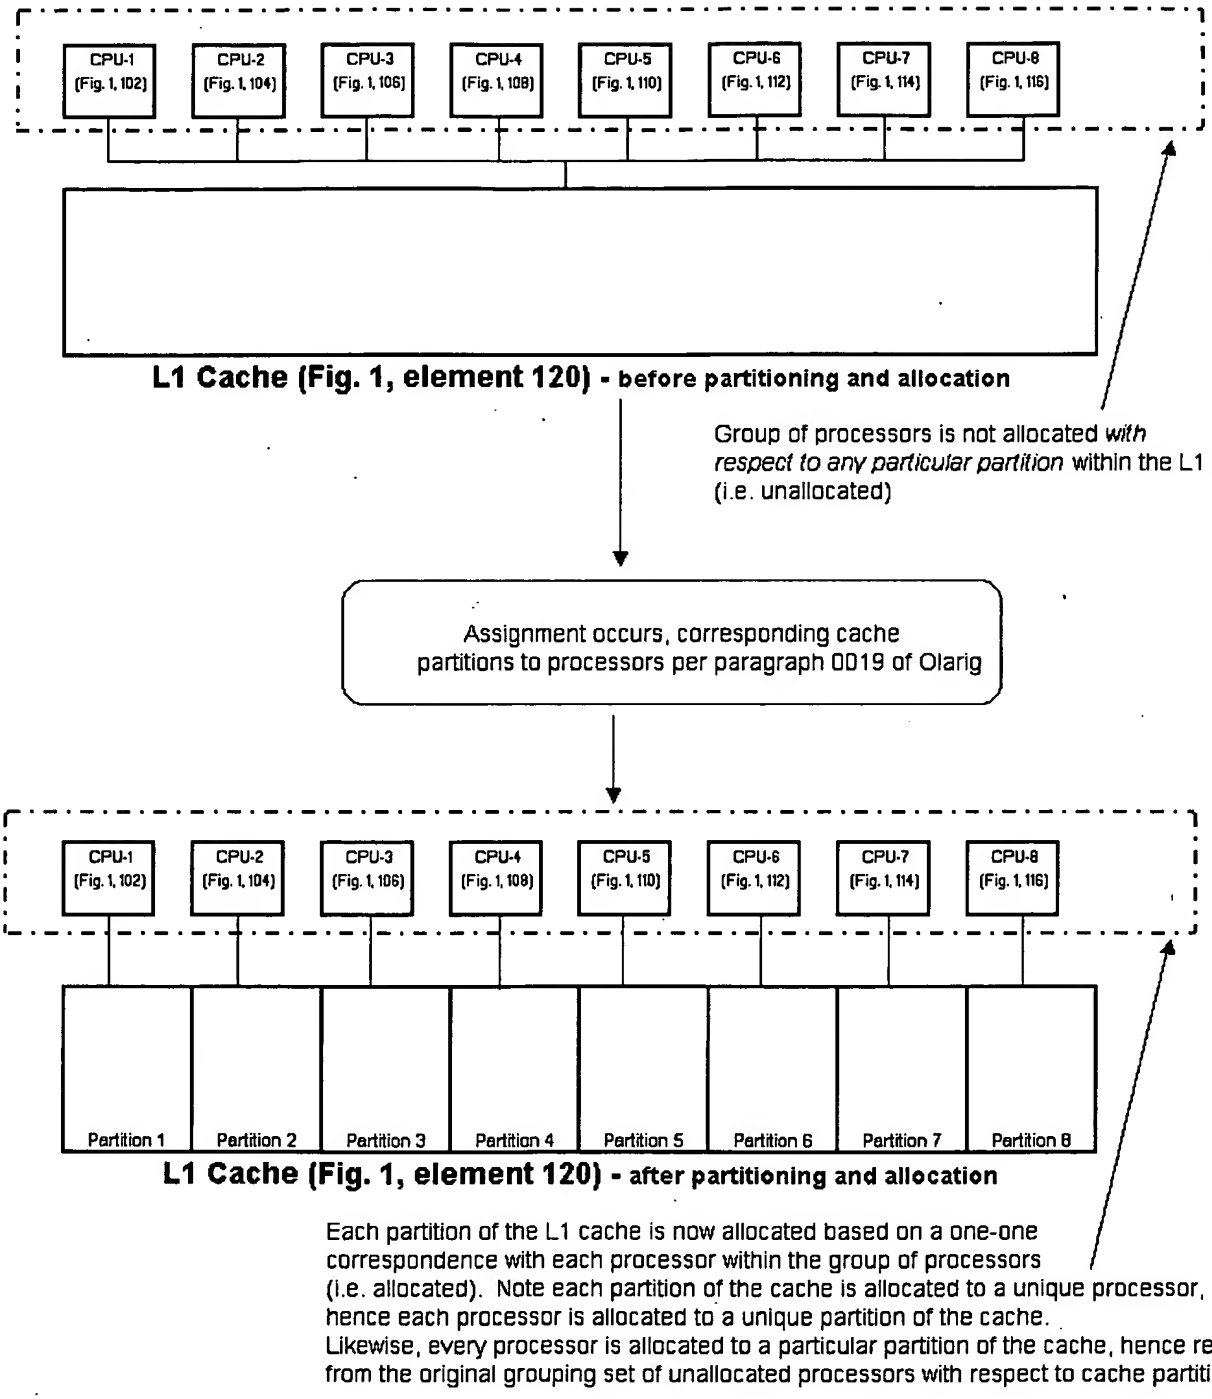

Examiner's Fig. 1

(all reference numerals correspond to Olarig, Fig. 1)

(11) Related Proceeding(s) Appendix

No decision rendered by a court or the Board is identified by the examiner in the Related Appeals and Interferences section of this examiner's answer.

Respectfully submitted,

Craig E Walter

Examiner

AU 2188

Conferees:

Hyung Sough

Supervisory Patent Examiner

AU 2188

3-1-07

Lynne Browne

Appeal Practice Specialist

TQAS 2100

Lynne H. Browne
Appeal Specialist, TQAS
Technology Center 2100